Problem Sets

- Problem Set 3
  - Distributed Tuesday, 3/18.
  - Due Thursday, 4/3
- Problem Set 4
  - Distributed Tuesday, 4/1
  - Due Tuesday, 4/15.
- Probably a total of 5 problem sets.

E-M

- Reading:
  - Forsyth & Ponce 16.1, 16.2
  - Forsyth & Ponce 16.3 for challenge problem.
- Yair Weiss: *Motion Segmentation using EM – a short tutorial.*
  - *www*
  - Especially 1st 2 pages.

Examples of Perceptual Grouping

- Boundaries:
  - Find closed contours near edges that are smooth.
  - Gestalt laws: common form, good continuation, closure.
- Smooth Paths:
  - Find open, smooth paths in images.
  - Applications: road finding, intelligent scissors (click on points and follow boundary between them).
  - Gestalt laws: Good continuation, common form.

Examples of Perceptual Grouping

- Regions: Find regions with uniform properties, to help find objects/parts of interest.
  - Color
  - Intensity
  - Texture
  - Gestalt laws: common form, proximity.

Examples of Perceptual Grouping

- Useful features:
  - Example, straight lines. Can be used to find vanishing points.
  - Gestalt laws: collinearity, proximity.

Parametric Methods

- We discussed Ransac, Hough Transform.
- The have some limitations
  - Object must have few parameters.
  - Finds an answer, but is it the best answer?
  - Hard to say because problem definition a bit vague.

Expectation-Maximization (EM)

- Can handle more parameters.
- Uses probabilistic model of all the data.
  - Good: we are solving a well defined problem.
  - Bad: Often we don't have a good model of the data, especially noise.
  - Bad: Computations only feasible with pretty simple models, which can be unrealistic.
- Finds (locally) optimal solution.

E-M Definitions

- Models have parameters: u
  - Examples: line has slope/intercept; Gaussian has mean and variance.
- Data is what we know from image: y
  - Examples: Points we want to fit a line to.
- Assignment of data to models: z
  - Eg., which points came from line 1.
  - z(i,j) = 1 means data i came from model j.
- Data and assignments (y & z): x.

E-M Definitions

- Missing Data: We know y. Missing values are u and z.
- Mixture model: The data is a mixture of more than one model.

E-M Quick Overview

- We know data (y).
- Want to find assignments (z) and parameters (u).
- If we know y & u, we can find z more easily.
- If we know y & z, we can find u more easily.
- Algorithm:
  1. Guess u.
  2. Given current guess of u and y, find z. (E)
  3. Given current guess of z and y, find u. (M)
  4. Go to 2.

Example: Histograms

- Histogram gives 1D clustering problem.
- Constant regions + noise = Gaussians.
- Guess mean and variance of pixel intensities.
- Compute membership for each pixel.
- Compute means as weighted average.
- Compute variance as weighted sample variance.
- Details: *whiteboard;* Also, Matlab and Weiss.

More subtle points

- Guess must be reasonable, or we won't converge to anything reasonable.
  - Seems good to start with high variance.
- How do we stop.
  - When things don't change much.
  - Could look at parameters (u).
  - Or likelihood of data.

Overview again

- Break unknowns into pieces. If we know one piece, other is solvable.
- Guess piece 1.
- Solve for piece 2. Then solve for 1. ….
- Very useful strategy for solving problems.

Drawbacks

- Local optimum.
- Optimization: we take steps that make the solution better and better, and stop when next step doesn't improve.
- But, we don't try all possible steps.

Local maximum

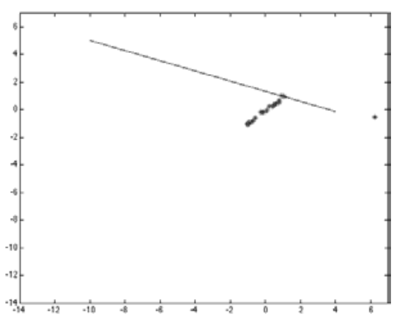

which is an excellent fit to some points

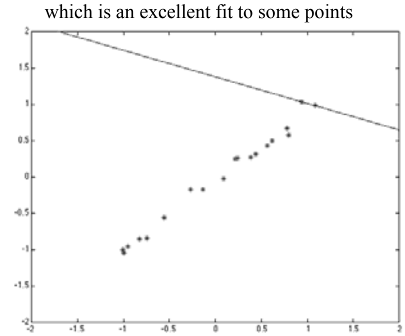

Drawbacks

- How do we get models?
  - But if we don't know models, in some sense we just don't understand the problem.
- Starting point? Use non-parametric method.
- How many models?

A dataset that is well fitted by four lines

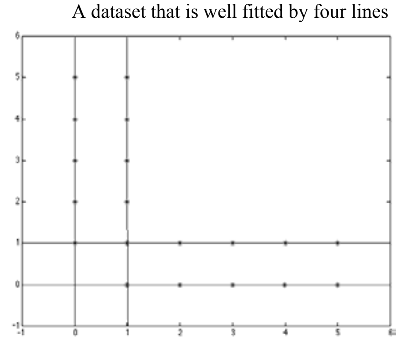

Result of EM fitting, with one line (or at least, one available local maximum).

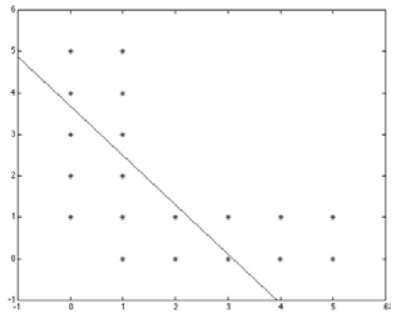

Result of EM fitting, with two lines (or at least, one available local maximum).

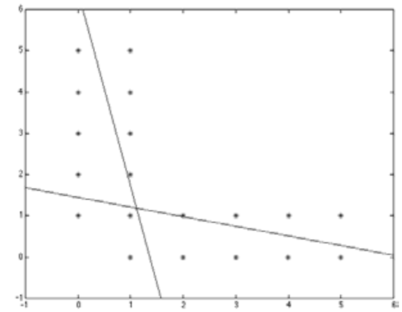

Seven lines can produce a rather logical answer

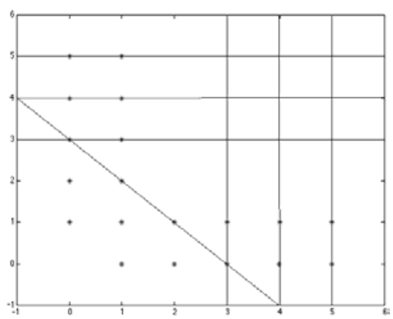

Segmentation with EM

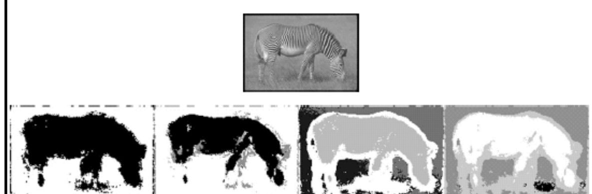

Figure from "Color and Texture Based Image Segmentation Using EM and Its Application to Content Based Image Retrieval",S.J. Belongie et al., Proc. Int. Conf. Computer Vision, 1998, c1998, IEEE

Motion segmentation with EM

- Model image pair (or video sequence) as consisting of regions of parametric motion
  - affine motion is popular

$$\begin{pmatrix} v_x \\ v_y \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix}\begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} t_x \\ t_y \end{pmatrix}$$

- Now we need to
  - determine which pixels belong to which region
  - estimate parameters

- Likelihood
  - assume

$$I(x,y,t) = I\left(x+v_x, y+v_y, t+1\right) + noise$$

- Straightforward missing variable problem, rest is calculation

Three frames from the MPEG "flower garden" sequence

Figure from "Representing Images with layers,", by J. Wang and E.H. Adelson, IEEE Transactions on Image Processing, 1994, c 1994, IEEE

Grey level shows region no. with highest probability

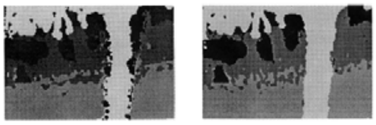

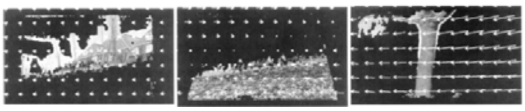

Segments and motion fields associated with them

Figure from “Representing Images with layers,”, by J. Wang and E.H. Adelson, IEEE Transactions on Image Processing, 1994, c 1994, IEEE

If we use multiple frames to estimate the appearance of a segment, we can fill in occlusions; so we can re-render the sequence with some segments removed.

Figure from “Representing Images with layers,”, by J. Wang and E.H. Adelson, IEEE Transactions on Image Processing, 1994, c 1994, IEEE

Probabilistic Interpretation

- We want P(u | y)
  - (A probability distribution of models given data).
- Or maybe P(u,z | y). Or argmax(u) P(u|y).
- We compute: argmax(u,z) P(y | u,z).
  - Find the model and assignments that make the data as likely to have occurred as possible.
  - This is similar to finding most likely model and assignments given data, ignoring prior on models.

Generalizations

- Multi-dimensional Gaussian.
  - Color, texture, …

$$\frac{1}{(2\pi)^{d/2}\det(\Sigma_i)^{1/2}}\exp(-\frac{1}{2}(\vec{x}-\vec{u})^T\Sigma^{-1}(\vec{x}-\vec{u}))$$

  - Examples: 1D reduces to Gaussian
  - 2D: nested ellipsoids of equal probability.
  - Discs if Covariance (Sigma) is diagonal.